

Amendments to the Claims

1. (Currently Amended) A general computer network controller, coupled to
2 ~~preferably operative in~~ a system area network, said controller comprising: including
a data buffer configured to handle one or more handling payloads; and
4 a fully associative context block configured to hold a plurality of last recently
used contexts to provide a dynamic resource allocation scheme reflecting run time
6 situations; and
a dedicated, programmable micro sequencer tightly coupled to said context block
8 and configured to: handling
control said context block; and
10 handle control flow and being capable of running different process
multiple types of network packets and protocols; being
12 wherein said micro sequencer is packet format independent and network
independent; and
14 ~~wherein said micro sequencer is tightly coupled to a fully associative context block for~~
~~control thereof, said context block being operative to hold a number of last recently used~~
16 ~~contexts to provide a dynamic resource allocation scheme reflecting run time situations,~~
wherein substantial parts of said contexts are being updated by said micro
18 sequencer, by an inbound scheduler and by a network protocol engine.

2. (Currently Amended) The computer network controller of claim 1, further
2 comprising: wherein said micro sequencer is operative to control
a scalable memory array configured ~~which can be used as a table for Inbound~~
4 address mapping of registered memory and access protection, and further configured as a
means for keeping context information about all active channels.

3. (Currently Amended) The computer network controller of claim 1,
2 wherein said fully associative context block couples ~~constitutes a connection between~~
said inbound scheduler and said network protocol engine, thereby facilitating an ability of
4 giving said network controller the ability to pipeline tasks and execute in parallel.

4. (Currently Amended) The computer network controller of claim 3,
2 wherein:
said context block is configured for dynamic allocation of ~~operative to have~~
4 contexts ~~dynamically allocated~~ between inbound remote direct memory access, inbound
remote memory access and outbound remote memory access;
6 ~~wherein two upper contexts are nevertheless being reserved for locally driven~~
remote direct memory access; and
8 said context block is configured to store ~~containing~~ information including one or
more of the following events:
10 - expected sequence number of ~~a~~ the next packet for sequence checking,
- input gathering size in order to optimize use of an attached bus,
A2 12 - packet type defined by the network for a specific virtual channel,
- accumulated message cyclic redundancy check for data integrity,
14 - source addresses,
- destination addresses,
16 - mapping for remote direct memory access operations,
- dedicated flags ~~like page crossing to~~ facilitate ~~do~~ new mapping,
18 - word count zero detection, and
- ~~as well as~~ protection tag check; and
20 wherein said ~~all these information events~~:
are received from said inbound scheduler, said micro sequencer and said
22 network protocol engine;
are to be ~~be~~ synchronized by said context block; and
24 are used by said micro sequencer to invoke, restart, switch or terminate a
thread immediately.

5. (Currently Amended) The computer network controller of claim 1,
2 wherein:
said micro sequencer is further configured ~~operative to~~ control said network
4 protocol engine; ~~which in its turn is operative~~

6 said network protocol engine is configured to perform link injection control,
based on feedback from a link layer and as well as intervention from an operative
system,; and
8 said network protocol engine is further configured ~~being operative~~ to schedule
packets to the network.

6. (Currently Amended) The computer network controller of claim 1,
2 wherein said inbound scheduler is configured ~~operative~~ to decode, schedule and invoke
running tasks or allocate new tasks, based on;

- 4 i) packets received from the network,
ii) memory mapped operations received from a bus attachment module,
6 iii) descriptors inserted in first-in, first-out work queues ~~fifos~~ by a user application,
and
8 iv) tasks received from said context block.

7. (Currently Amended) In a system area network comprising a plurality of
2 host channel adapters, a plurality of target channel adapters and a switching fabric, each
~~respective one of said adapter comprising: adapters being constituted by a computer~~
4 ~~network controller of the type defined in claim 1, together with a bus attachment module~~
~~and a network link interface;~~
6 a data buffer configured to handle one or more payloads;
a fully associative context block configured to hold a plurality of last recently
8 used contexts to provide a dynamic resource allocation scheme reflecting run time
situations; and
10 a dedicated, programmable micro sequencer tightly coupled to said context block
and configured to control said context block and handle control flow and process multiple
12 types of network packets and protocols;
a bus attachment module; and
14 a network link interface;
wherein said micro sequencer is packet format independent and network
16 independent, and wherein said contexts are updated by said micro sequencer, by an

inbound scheduler and by a network protocol engine,

18 a method for local and remote asynchronous completion control, the method
comprising:

20 detecting a final packet of a message directed from a local node to a remote node,
the final packet comprising:

22 an accumulated cyclic redundancy check covering the message; and

an address of a process completion queue on the remote node;

24 receiving the final packet at the remote node;

at the remote node:

26 performing an integrity check on the final packet;

signaling "receive complete" to the remote process completion queue; and

28 issuing a response to the final packet to the local node; and

at the local node, signaling "send complete" to a local process completion queue.

30 in which method as well accumulated message cyclic redundancy check as an

address to a remote completion queue, e.g. at a target, are attached, by a said micro

32 sequencer, to a last packet in a message to be sent from a sender, e.g. a host, to a receiver,
e.g. a target, whereby, on reception of said packet at said receiver and checking for data

34 integrity for the whole message by a target micro sequencer, "receive complete" is
signaled directly from said target micro sequencer in the remote process completion

36 queue, and simultaneously a response is made back to the sender, which will then signal
"send complete" and status directly to a local process.

8. (New) A protocol engine for a channel adapter configured to interface a
2 system area network with a network node, the protocol engine comprising:

4 an inbound scheduler configured to schedule one or more of the following for
each of a plurality of tasks: decoding, scheduling and invoking;

6 a multi-context micro sequencer configured to handle control flow for multiple
communication channels between the network node and the system area network,
wherein said multi-context micro sequencer is packet format independent and network
8 independent;

a context block configured to store a set of least recently used contexts, wherein

10 each said context corresponds to one of the communication channels;
a data buffer configured to buffer payloads of packets for the multiple
12 communication channels; and
a network protocol engine configured to schedule transmission of packets onto the
14 system area network.

9. (New) The protocol engine of claim 8, wherein said multi-context micro
2 sequencer is further configured to:
detect page boundary crossing and word count zero; and
4 perform an integrity check of a message, wherein the message comprises one or
more packets.

10. (New) The protocol engine of claim 8, wherein said multi-context micro
2 sequencer is further configured to perform integrated local and remote completion.

A2
11. (New) The protocol engine of claim 8, wherein a subset of said contexts
2 stored in said context block is reserved for outbound RDMA (Remote Direct Memory
Access).

12. (New) The protocol engine of claim 11, wherein the remainder of said
2 contexts in said set of contexts are dynamically allocated among inbound RDMA
(Remote Direct Memory Access), inbound RMA (Remote Memory Access) and
4 outbound RMA.

13. (New) The protocol engine of claim 8, wherein each said context stored in
2 said context block comprises one or more of:
a source address;
4 a destination address;
RDMA operation mapping;
6 expected sequence number of a next packet;
an accumulated cyclic redundancy check; and

8 a set of dedicated flags for performing one or more of:
word count zero detection;
10 packet integrity checking;
sequence error checking;
12 protection tag checking; and
data buffer management.

14. (New) The protocol engine of claim 8, wherein said data buffer comprises
2 a number of entries equivalent to the number of least recently used contexts stored in said
context block.

15. (New) The protocol engine of claim 8, wherein said data buffer
2 comprises:
multiple read ports; and
4 multiple write ports;
wherein said multiple read ports and multiple write ports facilitate processing of
6 multiple tasks in parallel by the protocol engine.

16. (New) The protocol engine of claim 8, further comprising:
2 one or more work queues configured to store descriptors inserted by applications
executing on the network node; and
4 an inbound scheduler configured to schedule processing of said descriptors.

17. (New) The protocol engine of claim 16, wherein said inbound scheduler is
2 further configured to schedule:
receipt of a packet from the system area network;
4 a memory-mapped operation received from the network node; and
a task received from said context block.

18. (New) The protocol engine of claim 8, further comprising:
2 a first connection coupling the protocol engine to an internal bus of the network

node; and

4 a second connection coupling the protocol engine to the system area network.

19. (New) The protocol engine of claim 18, further comprising:

2 a third connection coupling the protocol engine to an address translation table;
wherein the address translation table is configured to:

4 maintain inbound address mapping; and

store context information not currently stored in said context block.

20. (New) The protocol engine of claim 18, wherein the size of packets

2 exchanged between the protocol engine and the network node differ from the size of
packets exchanged between the protocol engine and the system area network.

A2